

What is claimed is:

1. A method of reorganizing a tablespace in a database comprising:

reading a row of data from the database;

analyzing the row of data read out from the database;

5 determining whether to eliminate or retain the row of data based on at least one predetermined rule;

reloading the row of data into the database when it is determined that the row of data complies with the at least one predetermined rule;

10 eliminating the row of data when it is determined that the row of data does not comply with the at least one predetermined rule;

rebuilding an index related to the database to include keys that correspond to the reloaded row of data; and

repeating the determining, reloading, eliminating and rebuilding for each row of data in the database.

15 2. The method of claim 1, wherein the at least one predetermined rule comprises a constraint.

20 3. The method of claim 2, wherein the constraint is a rule that restricts values in the database.

4. The method of claim 1, wherein the at least one predetermined rule is a requirement for referential integrity.

5. The method of claim 4, wherein the requirement for referential integrity is a rule requiring all non-null foreign keys in the database to correspond to an actual key in another database.

6. The method of claim 1, wherein the reloading step comprises:
reloading the row of data into the database in a first open row of the database, wherein the first open row is a first row in the database that contains no data.

7. The method of claim 1, wherein the reloading step comprises:
loading the row of data into a second database in a first open row of the second database, wherein
the first open row is a first row in the second database that contains no data.

8. The method of claim 1, wherein the row of data is eliminated by deleting the row of data.

9. A system for reorganizing a tablespace in a database comprising:
a reading device adapted to read out a row of data from the database;
an analyzing device adapted to analyze the row of data to determine whether to

eliminate or retain the row of data based on at least one predetermined rule;

a reloading device adapted to reload the row of data when it is determined that the row of data satisfies the at least one predetermined rule;

an eliminating device adapted to eliminate the row of data when it is determined that the row of data does not comply with the at least one predetermined rule; and

5 a rebuilding device adapted to rebuild an index related to the database to include keys that correspond to the reloaded row of data; wherein

each row of data in the database is read out from the database and analyzed.

10 10. The system of claim 9, wherein the at least one predetermined rule comprises a constraint.

11. The system of claim 10, wherein the constraint is a rule that restricts values in the database.

15 12. The system of claim 9, wherein the at least one predetermined rule is a requirement for referential integrity.

20 13. The system of claim 12, wherein the requirement for referential integrity is a rule requiring all non-null foreign keys in the database to correspond to an actual key in another database.

14. The system of claim 9, wherein the reloading device reloads the row of data into the database in a first open row of the database, wherein the first open row is a first row in the database that contains no data.

15. The system of claim 9, wherein the reloading device loads the row of data into a second database in a first open row of the second database, wherein the first open row is a first row in the second database that contains no data.

16. The system of claim 9, wherein the eliminating device eliminates the row of data by deleting the row of data.

17. A computer recording medium including computer executable code for reorganizing a tablespace in a database, said computer executable code including:

- reading code for reading out a row of the data from the database;
- analyzing code for analyzing the row of data read out from the database;
- determining code for determining whether to eliminate or retain a row of data based on at least one predetermined rule;
- reloading code for reloading the row of data when it is determined that the row of data complies with the at least one predetermined rule;
- eliminating code for eliminating the row of data when it is determined that the row of data does not comply with the at least one predetermined rule;
- rebuilding code for rebuilding an index related to the database to include keys

corresponding to the reloaded row of data; and

repeating code for repeating the determining, reloading, eliminating and rebuilding for each row of data in the database.

18. The computer recording medium of claim 17, wherein the at least one
5 predetermined rule comprises a constraint.

19. The computer recording medium of claim 18, wherein the constraint is a rule that restricts values in the database.

10 20. The computer recording medium of claim 17, wherein the at least one predetermined rule is a requirement for referential integrity.

21. The computer recording medium of claim 20, wherein the requirement for referential integrity is a rule requiring all non-null foreign keys in the database to correspond
15 to an actual key in another database .

22. The computer recording medium of claim 17, wherein the reloading code reloads the row of data into the database in a first open row of the database, wherein
the first open row is a first row in the database that contains no data.

20 23. The computer recording medium of claim 17, wherein the reloading code loads

the row of data into a second database in a first open row of the second database, wherein
the first open row is a first row in the second database that contains no data.

24. The computer recording medium of claim 17, wherein the eliminating code
comprises deleting code for deleting the row of data.

5

25. A method of reorganizing a tablespace in a database comprising the steps of:
partitioning the database and a related index into a plurality of partitions;
selecting one partition of the plurality of partitions of the database and a related
partition of the related index;

10

reading a row of data out from the selected one partition of the database;

analyzing the row of data read out from the selected one partition;

determining whether to eliminate or retain the row of data based on at least one
predetermined rule;

15

reloading the row of data into the selected one partition of the database when it is
determined that the row of data complies with the at least one predetermined rule;

eliminating the row of data when it is determined that the row of data does not comply
with the at least one predetermined rule;

rebuilding the related partition of the related index to include keys corresponding to
the reloaded row of data;

20

repeating the determining, reloading, eliminating and rebuilding for each row in the
selected one partition of the database; and

repeating the selecting, reading, analyzing, determining, reloading, eliminating and rebuilding for each partition in the database.

26. The method of claim 25, wherein the at least one predetermined rule comprises a constraint.

5

27. The method of claim 26, wherein the constraint is a rule that restricts values in the database.

28. The method of claim 25, wherein the at least one predetermined rule is a requirement for referential integrity.

10

29. The method of claim 28, wherein the requirement for referential integrity is a rule requiring all non-null foreign keys in the database to correspond to an actual key in another database.

15

30. The method of claim 25, wherein the reloading step comprises:
reloading the row of data into the selected one partition of the database in a first open row of the selected one partition of the database, wherein
the first open row is a first row in the selected one partition of the database that
contains no data.

20

31. The method of claim 25, wherein the reloading step comprises:
loading the row of data into a partition of a second database in a first open row of the
second database, wherein
the first open row is a first row in the second database that contains no data.

5 32. The method of claim 25, wherein the row of data is eliminated by deleting the
row of data.

33. A system for reorganizing a tablespace in a database comprising:
a partitioning device adapted to partition the database and a related index into a
10 plurality of partitions;
a partition selecting device adapted to select one partition of the plurality of partitions
of the database and a related partition of the related index;
a reading device adapted to read a row of data out from the selected one partition of
the database;
15 an analyzing device adapted to analyze the row of data read out from the selected one
partition and determine whether to eliminate or retain the row of data based on at least one
predetermined rule;
a reloading device adapted to reload the row of data into the selected one partition of
the database when it is determined that the row of data complies with the at least one
20 predetermined rule;
an eliminating device adapted to eliminate the row of data when it is determined that

the row of data does not comply with the at least one predetermined rule; and
a rebuilding device adapted to rebuilt the related partition of the related index to
include keys corresponding to the reloaded row of data; wherein
each row of the selected one partition of the database and each partition of the
plurality of partitions of the database are reorganized by the system.

5

34. The system of claim 33, wherein the at least one predetermined rule comprises a
constraint;

35. The system of claim 34, wherein the constraint is a rule that restricts values in the
10 database.

36. The system of claim 33, wherein the at least one predetermined rule is a
requirement for referential integrity.

15 37. The system of claim 36, wherein the requirement for referential integrity is a rule
requiring all non-null foreign keys in the database to correspond to an actual key in another
database.

38. The system of claim 33, wherein the reloading device reloads the row of data into
20 the selected one partition of the database in a first open row of the database, wherein
the first open row is a first row in the selected one partition of the database that

contains no data.

39. The system of claim 33, wherein the reloading device loads the row of data into a partition of a second database in a first open row of the second database, wherein the first open row is a first row in the second database that contains no data.

5

40. The system of claim 33, wherein the eliminating device eliminates the row of data by deleting the row of data.

41. A computer recording medium including computer executable code for reorganizing a tablespace in a database, said computer executable code including:

10 partitioning code for partitioning the database and a related index into a plurality of partitions;

partition selecting code for selecting one partition of the plurality of partitions of the database and a related partition of the related index;

15 reading code for reading a row of data out from the selected one partition;

analyzing code for analyzing the row of data read out from the selected one partition;

determining code for determining whether to eliminate or retain the row of data based on at least one predetermined rule;

reloading code for reloading the row of data into the selected one partition of the

20 database when it is determined that the row of data complies with the at least one predetermined rule;

eliminating code for eliminating the row of data when it is determined that the row of data does not comply with the at least one predetermined rule;

rebuilding code for rebuilding the related partition of the related index to include keys corresponding to the reloaded row of data;

row repeating code for repeating the determining, reloading, eliminating and
5 rebuilding for each row in the selected one partition of the database; and

partition repeating code for repeating the selecting, reading, analyzing, determining reloading, eliminating and rebuilding for each partition in the database.

42. The computer recording medium of claim 41, wherein the at least one
10 predetermined rule comprises a constraint.

43. The computer recoding medium of claim 42, wherein the constraint is a rule that restricts values in the database.

15 44. The computer recording medium of claim 41, wherein the at least one predetermined rule is a requirement for referential integrity.

45. The computer recording medium of claim 44, wherein the requirement for referential integrity is a rule requiring all non-null foreign keys in the database to correspond
20 to an actual key in another database .

46. The computer recording medium of claim 41, wherein the reloading code reloads the row of data into the selected one partition of the database in a first open row of the selected one partition of the database, wherein

the first open row is a first row in the selected one partition of the database that contains no data.

5

47. The computer recording medium of claim 41, wherein the reloading code loads the row of data into a partition of a second database in a first open row of the second database, wherein

the first open row is a first row in the second database that contains no data.

10

48. The computer recording medium of claim 41, wherein the eliminating code comprises deleting code for deleting the row of data.

49. A two-step method of reorganizing a tablespace in a database, the two-step method being repeated for each row in the database, comprising:

15

a first step including:

reading a row of data from the database;

analyzing the row of data read out from the database;

determining whether to eliminate or retain the row of data based on at least one

20 predetermined rule; and

a second step including:

reloading the row of data into the database when it is determined that the row of data
complies with the at least one predetermined rule;

eliminating the row of data when it is determined that the row of data does not comply
with the at least one predetermined rule; and

rebuilding an index related to the database to include keys that correspond to the
5 reloaded row of data.

50. A system for performing a two-step process of reorganizing a tablespace in a
database, the two-step process being repeated for each row in the database, the system
comprising:

10 a database; and

a processor for performing the two-step process comprising,

a first step including:

reading a row of data from the database;

analyzing the row of data read out from the database;

15 determining whether to eliminate or retain the row of data based on at least one
predetermined rule; and

a second step including:

reloading the row of data into the database when it is determined that the row of data
complies with the at least one predetermined rule;

20 eliminating the row of data when it is determined that the row of data does not comply
with the at least one predetermined rule; and

rebuilding an index related to the database to include keys that correspond to the reloaded row of data.

51. A computer recording medium including computer executable code for reorganizing a tablespace in a database in two steps, the two steps being repeated for each row of data in the database, the computer executable code comprising:

5 first step code including:

reading code for reading out a row of data from the database;

analyzing code for analyzing the row of data read out from the database;

determining code for determining whether to eliminate or retain the row of data based

10 on at least one predetermined rule; and

second step code including:

reloading code for reloading the row of data when it is determined that the row of data complies with the at least one predetermined rule;

eliminating code for eliminating the row of data when it is determined that the row of

15 data does not comply with the at least one predetermined rule; and

rebuilding code for rebuilding an index related to the database to include keys corresponding to the reloaded row of data.

52. A two-step method of reorganizing a tablespace in a database, the two-step method being repeated for each row in a partition of the database and each partition of a plurality of partitions of the database, the method comprising:

20

a first step including:

partitioning the database and a related index into the plurality of partitions:

selecting one partition of the plurality of partitions of the database and a related
partition of the related index;

reading a row of data out from the selected one partition of the database;

5 analyzing the row of data read out from the selected one partition;

determining whether to eliminate or retain the row of data based on at least one
predetermined rule; and

a second step including:

reloading the row of data into the selected one partition of the database when it is

10 determined that the row of data complies with the at least one predetermined rule;

eliminating the row of data when it is determined that the row of data does not comply
with the at least one predetermined rule; and

rebuilding the related partition of the related index to include keys corresponding to
the reloaded row of data.

15

53. A system for performing a two-step process of reorganizing a tablespace of a
database, the two-step process being repeated for each row in a partition of the database and
each partition of a plurality of partitions of the database, the system comprising:

a database; and

20 a processor for performing the two-step process comprising,

a first step including:

partitioning the database and a related index into the plurality of partitions;
selecting one partition of the plurality of partitions of the database and a related
partition of the related index;

reading a row of data out from the selected one partition of the database;

analyzing the row of data read out from the selected one partition;

5 determining whether to eliminate or retain the row of data based on at least one
predetermined rule; and

a second step including:

reloading the row of data into the selected one partition of the database when it is
determined that the row of data complies with the at least one predetermined rule;

10 eliminating the row of data when it is determined that the row of data does not comply
with the at least one predetermined rule; and

rebuilding the related partition of the related index to include keys corresponding to
the reloaded row of data.

15 54. A computer recording medium including computer executable code for
reorganizing a tablespace of a database in two steps, the two steps being repeated for each
row of data of a partition of the database and each partition of a plurality of partitions of the
database, the computer executable code comprising:

first step code including:

20 partitioning code for partitioning the database and a related index into the plurality of
partitions;

partition selecting code for selecting one partition of the plurality of partitions of the database and a related partition of the related index;

reading code for reading a row of data out from the selected one partition;

analyzing code for analyzing the row of data read out from the selected one partition;

and

5 determining code for determining whether to eliminate or retain the row of data based on at least one predetermined rule; and

second step code including:

reloading code for reloading the row of data into the selected one partition of the database when it is determined that the row of data complies with the at least one

10 predetermined rule;

eliminating code for eliminating the row of data when it is determined that the row of data does not comply with the at least one predetermined rule; and

rebuilding code for rebuilding the related partition of the related index to include keys corresponding to the reloaded row of data.

15

20